

PCT

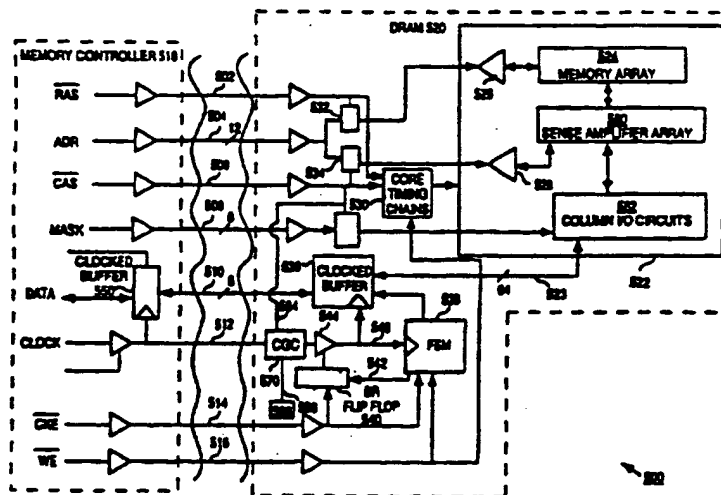
WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : <b>G06F</b>		A2	(11) International Publication Number: <b>WO 97/42557</b>
			(43) International Publication Date: 13 November 1997 (13.11.97)
(21) International Application Number: <b>PCT/US97/07463</b>		(81) Designated States: AL, AM, AT, AT (Utility model), AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, CZ (Utility model), DE, DE (Utility model), DK, DK (Utility model), EE, EE (Utility model), ES, FI, FI (Utility model), GB, GE, GH, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (Utility model), TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ARIPO patent (GH, KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: <b>1 May 1997 (01.05.97)</b>			
(30) Priority Data: <b>08/648,300 7 May 1996 (07.05.96) US</b>			
(71) Applicant: <b>RAMBUS, INC. [US/US]; 2465 Latham Street, Mountain View, CA 94040 (US).</b>			
(72) Inventors: <b>BARTH, Richard, M.; 787 Los Robles, Palo Alto, CA 94306 (US). HOROWITZ, Mark, A.; 2024 Columbia Street, Palo Alto, CA 94306 (US). HAMPEL, Craig, E.; 5927 Dunn Avenue, San Jose, CA 95123 (US). WARE, Frederick, A.; 13961 Fremont Pines Lane, Los Altos Hills, CA 94022 (US).</b>			
(74) Agents: <b>VINCENT, Lester, J. et al.; Blakely, Sokoloff, Taylor &amp; Zafman, 7th floor, 12400 Wilshire Boulevard, Los Angeles, CA 90025-1026 (US).</b>		Published <i>Without international search report and to be republished upon receipt of that report.</i>	

(54) Title: ASYNCHRONOUS REQUEST/SYNCHRONOUS DATA DYNAMIC RANDOM ACCESS MEMORY



(57) Abstract

A method and system for transferring information within a computer system is provided. The system includes a memory device that has a lower power mode in which data transfer circuitry is not driven by a clock signal, and a higher power mode in which data transfer circuitry is driven by a clock signal. The system further includes a memory controller that sends control signals to the memory device to initiate a data transfer transaction. The memory device receives the control signals asynchronously, and assumes the second mode in response to one of the control signals. While the memory device is in the second mode, the memory controller sends a control signal to identify a particular clock cycle. The memory device synchronously transfers the data. The memory device determines when to begin the data transfer based on the identified clock cycle and the type of data transfer that has been specified.

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SE	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

-1-

**ASYNCHRONOUS REQUEST/SYNCHRONOUS  
DATA DYNAMIC RANDOM ACCESS MEMORY**

**FIELD OF THE INVENTION**

The present invention relates to digital storage devices, and more specifically, to dynamic random access memory.

**BACKGROUND OF THE INVENTION**

Improvements in fabrication technology have resulted in dynamic random access memories (DRAMs) with increased density, faster performance, and higher operating frequencies. Because overall memory bandwidth requirements are rising and the number of DRAMs in a system is falling, the ability to quickly transport data to and from each DRAM has become increasingly important.

**ASYNCHRONOUS DRAMS**

In conventional memory systems, the communication between a memory controller and DRAMs is performed through asynchronous communications. For example, the memory controller uses control signals to indicate to the DRAM when requests for data transactions are sent. The data transfers themselves are also performed asynchronously. To meet increased speed requirements, various enhanced asynchronous memory systems have been developed. One such system is the Extended Data Out (EDO) DRAM memory system.

Figure 1 is a block diagram illustrating a typical EDO DRAM system 100. In the EDO DRAM system 100, data transfers are performed asynchronously in response to control signals and addresses sent from pin buffers 116 of a memory controller to pin buffers 118 of the EDO DRAM over a plurality of lines 120, 122, 124, 134 and 136. Specifically, lines 122 carry an address that is stored in latches 112 and 114. Line 120 carries a row address strobe ( $\overline{\text{RAS}}$ ) that controls when the address stored in latch 112 is sent to row decoder 106. Line 134 carries an output enable signal that controls data output of the DRAM. Line 136 carries a write enable signal that controls timing chains 108 and the direction of data flow on the bi-directional data bus 126.

Upon receiving an address, row decoder 106 loads data that corresponds to the address from a memory array 110 in memory core 102 into a sense amplifier array 130. Line 124 carries a column address strobe ( $\overline{\text{CAS}}$ ) that controls when the address stored in latch 114 is sent to column decoder 104. For a read

-2-

operation, the column decoder 104 causes the data that is stored in the columns of the sense amplifier array 130 that correspond to the address received by column decoder 104 to be transferred through column I/O circuits 132. The data passes through the column I/O circuits 132 to the memory controller over a data bus 126.

Alternately, an EDO DRAM may use address transition detect circuitry to initiate the retrieval of data from the memory core, rather than the  $\overline{\text{CAS}}$  signal. Address transition detect circuitry is circuitry that monitors the address bus to detect transitions in the data that is being sent on the address bus. When a transition is detected, the EDO DRAM restarts the timing chains causing data corresponding to a new address to fall out of the column I/O circuits 132.

The communication between the EDO DRAM and the memory controller is asynchronous. Thus, the EDO DRAM is not driven by an external clock. Rather, timing chains 108 that are activated by the  $\overline{\text{RAS}}$  and  $\overline{\text{CAS}}$  control signals are used to control the timing of the data transfer. Because the core 102 is not driven unless activated by the  $\overline{\text{RAS}}$  and  $\overline{\text{CAS}}$  control signals, the core 102 does not consume energy unless a data transfer operation is taking place. Therefore, the EDO DRAM consumes less power than alternative architectures in which the interface is clocked even when no memory operation is being performed.

Figure 2 is a timing diagram for a read operation in EDO system 100. At time T0 the memory controller places on lines 122 an address that indicates the bank and row from which data is to be read. At time T1 the  $\overline{\text{RAS}}$  signal goes LOW causing the address to be sent from latch 112 to row decoder 106. In response, row decoder 106 causes the appropriate row of data to be transferred from memory array 110 to sense amplifier array 130.

At time T2 the memory controller places on lines 122 the address of the column from which data is to be read. At time T3 the  $\overline{\text{CAS}}$  signal goes LOW causing the address to be sent from latch 114 to column decoder 104. In response, column decoder 104 sends through column I/O circuits 132 data from the selected column of the row stored in sense amplifier array 130. Assuming that  $\overline{\text{WE}}$  is HIGH and  $\overline{\text{OE}}$  is LOW, the data will appear on data bus 126. The data on the data bus 126 takes some time to stabilize. To ensure an accurate reading, the memory controller does not read the data from the data bus until time T4.

The delay between the time at which the  $\overline{\text{RAS}}$  signal goes LOW to initiate a read operation and the time at which the data may be read from the data bus 126

-3-

is identified as  $t_{RAC}$ . The delay between the time at which the  $\overline{CAS}$  signal goes LOW for a read operation and the time at which the data may be read from the data bus 126 is identified as  $t_{CAC}$ . The delay between the time at which the column address is placed on the address bus and the time at which the data may be read from the data bus 126 is identified as  $t_{CAA}$ . In a typical EDO DRAM, exemplary times are  $t_{CAC} = 15\text{ns}$  and  $t_{CAA} = 30\text{ns}$ .

In one variation, the memory controller is allowed to have column address flow through. The memory controller therefore has until T3 (the fall of  $\overline{CAS}$ ), rather than until T2 (the transmission of the column address), to decide whether to perform a given transaction. In the exemplary times above, the memory controller would have 15ns more time to decide whether to perform a given transaction.

#### SYNCHRONOUS DRAMS

DRAMs built with an asynchronous RAS/CAS interface have difficulty meeting the high memory bandwidth demands of many current computer systems. As a result, synchronous interface standards have been proposed. These alternative interface standards include Synchronous DRAMs (SDRAMs). In contrast to the asynchronous interface of EDO DRAMS, SDRAM systems use a clock to synchronize the communication between the memory controller and the SDRAMs. Timing communication with a clock allows data to be placed on the DRAM output with more precise timing. In addition, the clock signal can be used for internal pipelining. These characteristics of synchronous communication results in higher possible transfer rates.

Figure 3 is a block diagram illustrating a conventional SDRAM system 300. In system 300, the memory controller includes a plurality of clocked buffers 304 and the SDRAM includes a plurality of clocked buffers 306. Data from control line 310 and an address bus 312 are received by a finite state machine 308 in the SDRAM. The output of the finite state machine 308 and the address data are sent to memory array 302 to initiate a data transfer operation.

Figure 4 is a timing diagram that illustrates the signals generated in system 300 during a read operation. At time T0 the memory controller places a read request on line 310 and an address on bus 312. At time T1 the SDRAM reads the information on lines 310 and 312. Between T1 and T2 the SDRAM retrieves the data located at the specified address from memory array 302. At time T2 the

-4-

SDRAM places data from the specified address on data bus 314. At time T3 the memory controller reads the data off the data bus 314.

Because system 300 is synchronous, various issues arise that do not arise in asynchronous systems. Specifically, the synchronous system has numerous pipeline stages. Unbalanced pipeline stages waste computational time. For example, if a shorter pipeline stage is fed by a longer pipeline stage, there will be some period of time in which the shorter pipeline stage remains idle after finishing its operation and before receiving the next set of data from the preceding pipeline stage. Similarly, if a short pipeline stage feeds a longer pipeline stage, the shorter pipeline stage must wait until the longer pipeline stage has completed before feeding the longer pipeline stage with new input.

Each stage in the pipeline must allow for the setup, clock transition, and clock-to-output time of the flip-flop that is dividing the stages. Typically the execution time of each step is not substantially larger than the sum of these overheads, so the latency is significantly increased by them. Further, the memory controller may be running from a clock of a different frequency and/or phase from the DRAM subsystem clock. Crossing the boundaries between these clocks requires a time proportional to the clock frequencies. In addition, the architecture must take into account jitter that occurs when various data queues are clocked.

In general, the synchronous nature of the SDRAM architecture gives SDRAMs higher transfer rates than EDO DRAMs. However, the higher rates are achieved at the expense of increased latency and power consumption. Specifically, the time required to clock control and address data through various pipeline stages increases the delay between when an address for a read operation is transmitted and when the data from the specified address is actually supplied by the SDRAM.

The increased overhead (OV) that results from the use of synchronous transfer rather than an asynchronous transfer can be expressed by the formula  $OV = SD + (T_{DC} - D_1) + (T_{DC} - D_3) + (T_{DC} - (D_2 \text{ MOD } T_{DC}))$ , SD is synchronization delay,  $T_{DC}$  is the time period of the DRAM clock,  $D_1$  is the delay due to controller-to-DRAM time of flight,  $D_2$  is the time to perform a CAS operation,  $D_3$  is the delay due to DRAM-to-controller time of flight, and  $(D_2 \text{ MOD } T_{DC})$  is the remainder of  $(D_2/T_{DC})$ . SD is typically equal to  $(T_{DC} + T_{CC})$ , where  $T_{CC}$  is the duration of the controller clock cycle. In a system in which the external clock is at 66Mhz and the

-5-

DRAM subsystem clock is at 83Mhz, typical values may be:  $T_{DC}$  is 12ns,  $T_{CC}$  is 15ns,  $D_1$  is 6ns,  $D_2$  is 35ns, and  $D_3$  is 6ns. Thus, a typical OV would be  $(15 + 12) + (12 - 6) + (12 - 6) + (12 - 11) = 40ns$ .

Further, systems that use SDRAMs typically consume more power than the systems that use EDO DRAMs because, when the clock is enabled, the SDRAM interface is clocked whether or not a data transfer operation is actually being performed. For example, under typical conditions SDRAMs in an idle state consume approximately two to ten times more energy than EDO DRAMs in an idle state. When the clock is disabled, the clock must be enabled before a data transfer operation can be performed. More specifically, the clock must be enabled before any address or control information can be sampled by the SDRAM. The time used to enable the clock signal further increases the delay between the time that data is desired and the time that the requested data is available.

#### SUMMARY AND OBJECTS OF THE INVENTION

One object of the invention is to provide a memory system with an improved balance between request-to-data latency, power consumption and bandwidth.

According to one aspect of the invention, a memory interface is provided that maintains the high-bandwidth of synchronous systems, while reducing the latency and power requirements of these systems. This is accomplished by using an asynchronous interface for the address and control information, and using a synchronous interface for fast data transport.

According to one aspect of the invention, a controller transmits control signals requesting a data transfer to a memory device. The memory device asynchronously receives the control signals and synchronously performs the requested data transfer.

The memory device has a first mode in which data transfer circuits within the memory device are not driven by an internal clock signal. The memory device has a second mode in which data transfer circuits within the memory device are driven by the internal clock signal.

The memory device asynchronously receives the control signals. If the memory device is in the first mode, the memory device may assume the second mode in response to one or more of the control signals from a memory controller.

-6-

While in the second mode, the memory device transfers data with the data transfer circuits while the data transfer circuits are being driven by the internal clock signal. The memory device is also able to asynchronously perform data transfers while the memory device is in the first mode.

The internal clock signal is generated from an external clock signal that may selectively pass through a delay lock loop within the memory device. The memory device may support higher clock frequencies when the external clock signal passes through the delay lock loop to drive the data transfer circuits during a data transfer. Energy may be saved by circumventing the delay lock loop and using an external clock signal with a relatively slower frequency to drive the data transfer circuits during a data transfer.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

Figure 1 is a block diagram illustrating a prior art EDO DRAM system;

Figure 2 is a timing diagram illustrating the timing of signals when a read operation is performed in the EDO DRAM system of Figure 1;

Figure 3 is a block diagram illustrating a prior art SDRAM system;

Figure 4 is a timing diagram illustrating the timing of signals when a read operation is performed in the SDRAM system of Figure 3;

Figure 5a is a block diagram illustrating a memory system according to an embodiment of the present invention;

Figure 5b is a block diagram illustrating the clock generation circuitry of Figure 5a in greater detail;

Figure 6 is a timing diagram illustrating the timing of signals when a read operation is performed in the memory system of Figure 5a;

Figure 7 is a timing diagram illustrating the timing of signals when a write operation is performed in the memory system of Figure 5a;

Figure 8 is a flow chart illustrating the steps performed by a DRAM during a data transfer in the memory system shown in Figure 5a;



-7-

Figure 9 is a timing diagram illustrating the timing of signals during a read transaction in a memory system in which the mask and address signals are multiplexed over the same set of lines;

Figure 10 is a timing diagram illustrating the timing of signals during a read transaction in a memory system in which the mask, address and data signals are multiplexed over the same lines;

Figure 11 is a timing diagram illustrating the timing of signals during a write transaction in a memory system in which the mask, address and data signals are multiplexed over the same lines;

Figure 12 is a diagram illustrating the correlation between a  $\overline{\text{RAS}}$  signal and separate PRECHARGE and SENSE signals according to an embodiment of the invention.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring to Figure 5a, it illustrates a memory system 500 according to one embodiment of the invention. System 500 includes a memory controller 518 coupled to a DRAM 520 by a plurality of lines. The lines connecting memory controller 518 to DRAM 520 include a  $\overline{\text{RAS}}$  line 502, an address bus 504, a  $\overline{\text{CAS}}$  line 506, a mask bus 508, a data bus 510, a clock line 512, a clock enable ( $\overline{\text{CKE}}$ ) line 514, and a write enable ( $\overline{\text{WE}}$ ) line 516.

The DRAM 520 has a core 522 that includes a memory array 524, a sense amplifier array 560, column I/O circuits 562, a row decoder 526 and a column decoder 528. DRAM 520 further includes core timing chains 530, latches 532 and 534, a clocked buffer 536, a finite state machine 538, a flip flop circuit 540 and a clock buffer 544. DRAM 520 further includes a clock generation circuit (CGC) 570 and a mode control circuit 566. As shall be explained in greater detail below, in system 500, data transfers are requested using asynchronous communication, while the actual data transfers performed in response to the requests may be performed using synchronous communication or asynchronous communication, depending on the mode selected by mode control circuit 566.

DRAM 520 illustrates an embodiment that supports at least three data transfer modes: synchronous mode, slow synchronous mode, and asynchronous mode. As shall be described in greater detail hereafter, mode control circuit 566 controls the transfer mode currently in effect in DRAM 520. Mode control circuit

566 may be, for example, a value stored in a register of DRAM 520, or pin coupled to an external control line.

### ASYNCHRONOUS TRANSFER CONTROL SIGNALS

Referring to Figure 6, it illustrates the timing of signals generated during an exemplary read transaction performed in memory system 500. Initially, the  $\overline{\text{CKE}}$  line 514 is HIGH, causing the flip flop 540 to turn off the clock buffer 544. When the clock buffer 544 is off, any clock signal on line 512 is prevented from driving the synchronous components of DRAM 520.

Prior to time T1, memory controller 518 transmits a row address over the address bus 504. The row address is loaded into latches 532 and 534. At time T1, memory controller 518 causes the  $\overline{\text{RAS}}$  signal to go LOW. When the  $\overline{\text{RAS}}$  line goes LOW, the row address passes through latch 532 and is applied to row decoder 526. For multi-bank devices, row and column addresses include bank address information. Row decoder 526 causes a sense operation to be performed on cells within memory array 524 to load into sense amplifier array 560 the data that corresponds to the specified row address.

Prior to time T3, the memory controller 518 places a column address on the address bus 504. The column address is loaded into latches 532 and 534. At time T3, the memory controller 518 causes the signal on the  $\overline{\text{CAS}}$  line 506 to go LOW. When the  $\overline{\text{CAS}}$  signal goes LOW, the column address passes through latch 534 and is applied to column decoder 528. Column decoder 528 causes an access operation to be performed on the data currently stored in the sense amplifier array 560. The data corresponding to the specified column address is sent from sense amplifier array 560 through column I/O circuits 562.

Prior to time T6, the memory controller 518 places a second column address on the address bus 504. The column address is loaded into latches 532 and 534. At time T6, the memory controller 518 causes the signal on the  $\overline{\text{CAS}}$  line 506 to go LOW. When the  $\overline{\text{CAS}}$  signal goes LOW, the column address passes through latch 534 and is applied to column decoder 528. Column decoder 528 causes an access operation to be performed on the data currently stored in the sense amplifier array 560. The data corresponding to the specified column address is sent from sense amplifier array 560 through column I/O circuits 562.

-9-

Significantly, all of the control information to perform the read transaction is sent without reference to any particular clock cycles. This is true even though the internal clock of DRAM 520 may be active at the time at which some of the control signals are sent by the memory controller 518. For example, the internal clock is active at time T6 when the second column address is latched to the column decoder. However, the timing of the  $\overline{\text{CAS}}$  signal that causes the column address to be latched and sent to the column decoder does not necessarily have any correlation with the clock signal.

According to an alternate embodiment, the  $\overline{\text{RAS}}$  and  $\overline{\text{CAS}}$  signals are well controlled with respect to the clock signal on line 512. In such an embodiment, the request for data is issued from the same time domain as the data transport clock on line 512. The timing of the data transfer may then be determined based on the timing of the control signals.

For example, data transport may begin a predetermined number of clock cycles after the clock cycle on which the falling edge of  $\overline{\text{CAS}}$  occurs. Various techniques may be used to ensure that the DRAM is aware of the clock cycle on which the falling edge of  $\overline{\text{CAS}}$  occurred. For example, the memory controller 518 may cause the  $\overline{\text{CKE}}$  signal to go LOW sufficiently before the falling edge of  $\overline{\text{CAS}}$  to ensure that the internal clock signal is stabilized by the time the falling edge of  $\overline{\text{CAS}}$  occurs.

Alternately, if it has been determined exactly how many clock cycles will elapse between the falling edge of  $\overline{\text{CKE}}$  and the first clock cycle to cause data to be output, the falling edge of  $\overline{\text{CKE}}$  may be used to time the data transfer. Under these circumstances, the internal clock signal must be relatively stable as soon as clock buffer 544 is activated in response to the falling edge of  $\overline{\text{CKE}}$ .

In an alternate embodiment, DRAM 520 includes address transition detect circuitry, thus avoiding the need of separate  $\overline{\text{CAS}}$  line 506 and signal. Address transition detect circuitry monitors the address bus 504 to detect transitions in the data that is being sent on the address bus 504. When a transition is detected, DRAM 520 restarts the timing chains 530 causing data corresponding to a new address to fall out of the column I/O circuits 562.

-10-

### MULTI-MODE CLOCK GENERATION CIRCUIT

As mentioned above, DRAM 520 supports a synchronous data transfer mode, a slow synchronous data transfer mode, and an asynchronous data transfer mode. The current data transfer mode determines the signal at the output of clock generation circuit 570. The mode is selected by mode control circuit 566, which may be, for example, a value in a register within DRAM 520 or a pin connected to an external control line.

Referring to Figure 5b, it is a block diagram illustrating a clock generation circuit 570 according to one embodiment of the invention. Clock generation circuit 570 includes a multiplexer 574 with three inputs and one output. The inputs are coupled to lines 564, 576 and 512. The output of multiplexer 574 feeds the input of clock buffer 544. The multiplexer 574 is controlled by the control signal sent by the mode control circuit 566 over line 568.

When mode control circuit 566 applies a signal associated with the asynchronous transfer mode over line 568, the signal on  $\overline{\text{CAS}}$  line 506 passes from line 564 through multiplexer 574 to the input of clock buffer 544. When mode control circuit 566 applies a signal associated with the synchronous transfer mode over line 568, the output of a delay lock loop (DLL) 572 whose input is the external clock signal on line 512 passes through multiplexer 574 to the input of clock buffer 544. When mode control circuit 566 applies a signal associated with the slow synchronous transfer mode over line 568, the external clock signal on line 512 passes through multiplexer 574 to the input of clock buffer 544.

### CLOCK ACTIVATION

Returning again to the transaction illustrated in Figure 6, the memory controller 518 causes the signal on the  $\overline{\text{CKE}}$  line 514 to go LOW at time T2. When the  $\overline{\text{CKE}}$  line 514 goes LOW, the flip flop 540 turns on the clock buffer 544. The signal at the output of clock generation circuit 570 begins driving clocked buffer 536 and finite state machine 538 through the clock buffer 544 when the clock buffer 544 is on. Assuming that DRAM 520 is in either synchronous mode or slow synchronous mode, the signal at the output of clock generation circuit 570 will reflect a clock signal on line 512, as shown in Figure 6.

In Figure 6, the "internal clock signal" is the signal generated at the output of the clock buffer 544. The internal clock signal generated by clock buffer 544

-11-

typically takes a few cycles to stabilize. Therefore, the signal on line 546 does not immediately reflect the clock signal on line 512. In the illustrated example, the signal on line 546 does not stabilize until some time has elapsed after T2.

In the illustrated read transaction, the source of the clock signal on line 512 is not activated until time T2. By turning off the source of the external clock signal when no data transfers are in progress, both the external clock source and DRAM 520 conserve power when data transfer operations are not being performed. In alternative embodiments, the source of the external clock signal on line 512 remains on, while the internal clock signal on line 546 is only turned on when DRAM 520 is actually involved in a data transfer, as described above.

### SYNCHRONIZING THE TIMING

After time T3, the DRAM 520 has all the information it requires to transmit data from the specified row and column, but does not yet know when to begin sending the data. In conventional SDRAMs, the timing of the data transfer is based on the timing of the data transfer request. Thus, if a controller sends a data transfer request on a particular clock cycle, then the controller knows that the requested data transfer will begin a predetermined number of clock cycles after the particular clock cycle.

In system 500, the data transfer requests are transmitted in an asynchronous manner. In fact, the clock source whose signal is used to time synchronous data transfers may not even be active at the time the  $\overline{\text{RAS}}$  and  $\overline{\text{CAS}}$  signals are transmitted. Therefore, the transmission of information other than the data itself (e.g.  $\overline{\text{CAS}}$ ,  $\overline{\text{RAS}}$ , address information,  $\overline{\text{WE}}$ , etc.) need not be associated with any particular clock cycle or mode. Consequently, DRAM 520 cannot time the transmission of data based on a clock cycle on which the  $\overline{\text{RAS}}$  or  $\overline{\text{CAS}}$  signals were transmitted, and memory controller 518 cannot use a clock cycle on which the  $\overline{\text{RAS}}$  or  $\overline{\text{CAS}}$  signals were transmitted to determine the clock cycle on which DRAM 520 will begin sending data.

According to one embodiment of the invention, the rising edge of the  $\overline{\text{CKE}}$  signal is used as a timing mark to indicate to the finite state machine 538 of DRAM 520 when to begin sending requested data. Specifically, the clock buffer 544 is activated at the falling edge of the  $\overline{\text{CKE}}$  signal (at time T2) as described above. The memory controller 518 continues to generate the LOW  $\overline{\text{CKE}}$  signal.

After the clock signal from the clock buffer has stabilized, the memory controller 518 causes the  $\overline{\text{CKE}}$  signal to go HIGH. The time at which the  $\overline{\text{CKE}}$  signal goes HIGH is used by memory controller 518 and the finite state machine 538 as a timing mark.

According to an alternate embodiment, a control line separate from the  $\overline{\text{CKE}}$  line 514 may be used to provide the timing mark. In an embodiment that uses a separate control line for the timing mark,  $\overline{\text{CKE}}$  might be full swing CMOS while the timing mark is low swing high speed signal.

In the illustrated read transaction, the first timing mark occurs at time T4. In one embodiment, the finite state machine 538 begins the transmission of the requested data a predetermined number of clock cycles after the timing mark, and memory controller 518 knows to expect the data from DRAM 520 the predetermined number of clock cycles after the timing mark. The predetermined number may be a fixed constant, or a value stored in a register within DRAM 520.

In the illustrated embodiment, the  $\overline{\text{WE}}$  signal is sampled at the first rising edge of the clock signal after  $\overline{\text{CKE}}$  is sampled HIGH. The sample of  $\overline{\text{WE}}$  is used to determine whether the transaction is going to be a read transaction or a write transaction. In the example shown in Figure 6,  $\overline{\text{WE}}$  is HIGH on the first rising edge of the clock signal after  $\overline{\text{CKE}}$  goes HIGH, indicating that the data transfer is going to be a read transaction.

In an alternate embodiment, the  $\overline{\text{WE}}$  signal can be sampled at the falling edge of  $\overline{\text{CKE}}$ . To increase transmission bandwidth, the input receive path of DRAM 520 and the output transmit path of DRAM 520 can be separately compensated. For example, clock generation circuit 570 and clock buffer 544 may be replaced with two clock generation circuit/clock buffer combinations, where one clock generation circuit/clock buffer combination is used to drive clocked buffer 536 to receive data and a different clock generation circuit/clock buffer combination is used to drive clocked buffer 536 to transmit data. Power is saved by activating only the clock generation circuit/clock buffer combination that will be involved in the transfer. By sampling the  $\overline{\text{WE}}$  signal at the falling edge of  $\overline{\text{CKE}}$ , the DRAM has more time between when the type of transaction (read or write) is known and when the data transfer will begin. During this interval the DRAM activates the clock buffer that corresponds to the type of transaction to be performed.

### THE SYNCHRONOUS DATA TRANSFER

As mentioned above, the finite state machine 538 causes data from the specified column of the specified row to be sent over the data bus 510 a predetermined number of clock cycles after the timing mark. The delay between a  $\overline{\text{CAS}}$  signal and the transmission of data must be long enough for the data from the appropriate column to be loaded through the column I/O circuits 562 into the clocked buffer 536. In the illustrated example, each column address corresponds to eight bytes. However, a packet size of eight bytes is merely exemplary. The actual size of data packets will vary from implementation to implementation. The present invention is not limited to any particular size of data packet.

In response to the  $\overline{\text{CAS}}$  signal, eight bytes that correspond to the specified column address are loaded through the column I/O circuits 562 into clocked buffer 536. During a data transfer, finite state machine 538 causes the eight bytes to be sent sequentially (per data bus width) from the clocked buffer 536 to the data bus 510. The clock signal from the clock buffer 544 determines the timing of the transmission of the eight bytes. In the illustrated example, two bytes are sent per clock cycle, beginning a time T5. The same clock signal is applied to a clocked buffer 550 in the memory controller 518. The eight bytes of data are sequentially received at the clocked buffer 550 based on the timing of the clock signal.

In the embodiment described above, data is transferred through column I/O circuits 562 to clocked buffer 536 eight bytes at a time, and transferred out of clocked buffer 536 to clocked buffer 550 one byte at a time. Consequently, in this embodiment, clocked buffer 536 may be a parallel to serial shift register, while clocked buffer 550 may be a serial to parallel shift register. The buffer circuits used to perform the transfer function may vary from implementation to implementation. The present invention is not limited to any particular type of clocked buffers, nor any particular clock speeds or bandwidths used to transfer data within DRAM 520 or between DRAM 520 and memory controller 518.

### SHUTDOWN AFTER A DATA TRANSFER

As mentioned above, the synchronous components within DRAM 520 begin to be driven at time T2 by the clock signal on line 512. While these components are being driven by the clock signal, the DRAM 520 continues to

-14-

consume relatively large amounts of power. DRAM 520 would continue to consume large amounts of power even when DRAM 520 is not involved in a data transfer if the DRAM 520 is not isolated from the clock signal on line 512 after the completion of a data transfer. Therefore, finite state machine 538 contains a mechanism for turning off clock buffer 544 after all of the outstanding data transfer operations that involve DRAM 520 have been completed.

According to one embodiment of the invention, finite state machine 538 uses a countdown timer to determine when to turn off the clock buffer 544. Specifically, upon detecting the timing mark, finite state machine 538 stores a count value in a countdown timer and begins decrementing the count value during each clock cycle. As shall be explained in greater detail below, the countdown timer is incremented or reloaded for each data block in multiple-block transfers. When count value of the countdown timer reaches zero, the finite state machine 538 sends a signal to flip flop 540 over a line 542 to cause the flip flop 540 to turn off the clock buffer 544. When clock buffer 544 is turned off, the synchronous components of DRAM 520 cease to be driven by the clock signal on line 512, causing DRAM 520 to assume a state in which little power is consumed.

Alternative embodiments may use other mechanisms for turning off the clock buffer 544 when all data transfers involving the DRAM 520 have been completed. For example, logic circuits within finite state machine 538 may be configured to detect the completion of a data transfer operation and determine whether there is any outstanding transaction that involves DRAM 520. If there is an outstanding transaction, then the finite state machine 538 transmits the appropriate signals to the clocked buffer 536 to perform the outstanding transaction. If there are no outstanding transactions that involve DRAM 520, then the finite state machine 538 sends a signal to the flip flop 540 to cause the clock buffer 544 to be turned off.

#### ASYNCHRONOUS DATA TRANSFER MODE

To achieve high data transfer rates, synchronous transfers can be performed as described above. However, under certain conditions it may be desirable to avoid the relatively high power consumption requirements of DLL 572 by performing data transfers asynchronously.



-15-

To perform asynchronous data transfers, mode control circuit 566 applies a control signal to line 568 to cause the signal on  $\overline{\text{CAS}}$  line 506 to be generated at the output of clock generation circuit 570, as described above. Memory controller 518 may then toggle the  $\overline{\text{CAS}}$  signal without introducing any new address information on address bus 504, causing the  $\overline{\text{CAS}}$  signal to act as a clock to drive clocked buffer 536.

In an alternate embodiment, asynchronous transfers may be performed by placing the clocked buffer 536 in flow-through. To address the width mismatch between the internal data bus 523 and the external data bus 510, the memory controller 518 presents sufficient addressing information to the DRAM 520 to select a single byte from the eight bytes loaded on the sense amplifier array 560.

#### SLOW SYNCHRONOUS DATA TRANSFER MODE

Even when synchronous data transfers are desired, the delay lock loop circuit 572 within clock generation circuit 570 may be bypassed to reduce power consumption. To bypass DLL 572, mode control circuit 566 applies a control signal to line 568 to cause the signal on line 512 to be generated at the output of clock generation circuit 570, as described above.

However, DRAM 520 cannot support the same transfer rate without the clock synchronization provided by the DLL 572 as it can with a clock synchronized by the DLL 572. Consequently, when the DLL 572 is bypassed, a slower clock signal must be used to perform the synchronous data transfers. Due to the lower clock frequency, the synchronous data transfers take longer than when the DLL 572 is used. Consequently, DRAM 520 is said to be in "slow" synchronous data transfer mode when an external clock signal that has not been phase compensated by a DLL is used to drive the data transfers.

#### MULTIPLE-BLOCK TRANSFERS

In an embodiment that uses a countdown timer, the seed count value used by the countdown timer is based on the amount of time required for DRAM 520 to send one packet of data (eight bytes in the illustrated example). Specifically, after the timing mark is detected, the clock buffer 544 must stay on long enough for a packet of data to be accessed, loaded, and transmitted from DRAM 520. If a new column address arrives before data from the previous column address has been

completely transferred, then the clock buffer 544 should stay on until the data from the new column address has been transmitted.

To prevent clock buffer 544 from being turned off between consecutive packet transfers, the finite state machine 538 adds a predetermined value to the count value in the countdown timer upon detecting a falling edge of the  $\overline{\text{CKE}}$  signal. Because the count value in the countdown timer is increased, a greater number of clock cycles will elapse before the count value reaches zero. Preferably, the predetermined value that is added to the count value causes the shutdown of the clock buffer 544 to be delayed long enough for the additional packet of data to be transferred.

In an alternate embodiment, a predetermined value is loaded into the countdown timer upon detecting a falling edge of the  $\overline{\text{CKE}}$  signal. The predetermined value is large enough to ensure that the countdown timer will not reach zero before a packet of data is transferred. During the transfer of multiple data packets the counter will repeatedly be reloaded and thus never reach zero.

Figure 6 illustrates the timing of a read transaction in which two data packets are transferred. At time T2 the count value in the countdown timer is set to a value that ensures that the clock buffer 544 will be on long enough for one packet of data to be transferred. In the illustrated example, the count value will be set to a value that ensures that the clock buffer 544 remains on until time T8.

At time T5, the finite state machine 538 adds or reloads a predetermined number to the count value in response to detecting the falling edge of the  $\overline{\text{CKE}}$  signal. At time T8 the transmission of the first packet of data is completed. Because the predetermined value was added to the count value or the counter was reloaded, the finite state machine 538 does not turn off the clock buffer at T8. Rather, the count value does not reach zero until after time T9, when the packet of data from the second column has been completely transferred.

According to an alternate embodiment of the invention, finite state machine 538 contains logic for keeping track of how many CAS requests remain to be serviced. Upon detecting the falling edge of the  $\overline{\text{CKE}}$  signal, the finite state machine 538 increments the outstanding request value. Upon completing the transfer of one data block, the finite state machine 538 decrements the outstanding request value. When the outstanding request value is zero, the finite state machine 538 turns off the clock buffer 544.

### WRITE TRANSACTION TIMING

Figure 7 is a timing diagram of the signals generated in system 500 during a two packet write transaction. At time T1, an address on address bus 504 is transferred from latch 532 to row decoder 526 when the  $\overline{\text{RAS}}$  signal on line 502 goes LOW. At time T2, the  $\overline{\text{CKE}}$  signal on line 514 goes LOW causing the flip flop 540 to activate the clock buffer 544. The finite state machine 538 detects the rising edge of the  $\overline{\text{CKE}}$  signal to determine that time T3 (the rising edge of the  $\overline{\text{CKE}}$  signal) is the timing mark for the first packet transfer.

The  $\overline{\text{WE}}$  signal is sampled on the first clock cycle after the rising edge of the  $\overline{\text{CKE}}$  signal to determine whether the transaction is going to be a read transaction or a write transaction. In the present example,  $\overline{\text{WE}}$  is LOW at the rising edge of the first clock cycle after the  $\overline{\text{CKE}}$  signal goes HIGH, indicating that the data transfer is going to be a write transaction.

At time T4,  $\overline{\text{CAS}}$  goes LOW indicating to DRAM 520 that the data bus 510 has data that is to be written to the column that corresponds to the address on the address bus 504. In the illustrated example, the address on the address bus at time T4 specifies column A. The DRAM also receives the mask data on lines 508. The finite state machine 538 controls the clocked buffer 536 to cause the data to be synchronously stored in column A in the sense amplifier array 560. Finite state machine 538 knows to expect the data at time T5 because time T5 is a predetermined number of clock cycles (e.g. one clock cycle) after the timing mark.

A timing mark occurs at time T6, the first clock cycle after another rising edge of the  $\overline{\text{CKE}}$  signal. At time T6 the  $\overline{\text{WE}}$  signal is sampled. In the illustrated example, the  $\overline{\text{WE}}$  signal is LOW at time T6 indicating a second column of data is to be written to DRAM 520. When receipt of the first packet of data is complete, the received data is stored in column A of the appropriate row of the memory array 524.

At time T7,  $\overline{\text{CAS}}$  goes LOW, indicating to DRAM 520 that the data bus 510 has data that is to be written to the column that correspond to the address on the address bus 504. In the illustrated example, the address on the address bus 504 at time T7 specifies column B. The finite state machine 538 controls the clocked buffer 536 to cause the data to be synchronously received into column B of the sense amplifier array 560. Finite state machine 538 knows to expect the

-18-

data at time T7 because time T7 is a predetermined number of clock cycles after the second timing mark at time T6. At time T8 the second packet of data has been completely received, so the DRAM 520 stores the second packet of data in column B of the appropriate row within memory array 524.

Figure 8 is a flowchart illustrating the operation of system 500 according to an embodiment of the invention. Initially, DRAM 520 is in a powered down state where clock buffer 544 is off. At step 800, the DRAM 520 detects the fall of the  $\overline{\text{RAS}}$  signal. The fall of the  $\overline{\text{RAS}}$  signal causes the address on the address bus 504 to be sampled (i.e. sent from latch 532 to row decoder 526) at step 802. At step 804 the core of DRAM 420 senses the row of data that corresponds to the address sampled at step 802.

At step 812, the DRAM 520 detects the falling edge of the  $\overline{\text{CAS}}$  signal. In response to detecting the falling edge of the  $\overline{\text{CAS}}$  signal, DRAM 520 samples the mask signals on lines 508 and the column address on the address bus 504 (step 814).

Steps 800, 802, 804, 812 and 814 are performed asynchronously and therefore do not require an active clock signal. Steps 806, 808 and 810 may occur before, in parallel with, or after steps 800, 802, 804, 812 and 814, and therefore are shown as a separate thread of execution.

At step 806, the  $\overline{\text{CKE}}$  signal goes LOW causing the clock buffer to be turned on at step 808. At step 810 the DRAM 520 detects a timing mark. In the embodiment described above, the timing mark is detected when the finite state machine 538 senses the start of the first clock cycle subsequent to the rising edge of the  $\overline{\text{CKE}}$  signal. The  $\overline{\text{WE}}$  signal is sampled at this time to determine whether the data transfer is going to be a read transaction or a write transaction.

At step 816, the finite state machine 538 determines the clock cycle on which the data transfer is to begin based on when the timing mark was detected and whether the transaction is a write transaction or a read transaction.

At step 818, it is determined whether the  $\overline{\text{WE}}$  signal sampled at step 810 indicated that the transaction is a write transaction. The  $\overline{\text{WE}}$  signal sampled at step 810 indicated that the transaction is a write transaction, then control proceeds to step 828. Otherwise, control proceeds to step 820.

Significantly, all of the steps performed up to step 816 are performed in an asynchronous manner. The use of an asynchronous mechanism to perform these

-19-

steps reduces the latency between the fall of the  $\overline{\text{RAS}}$  signal and the time that the appropriate row of data is sensed. By the time step 810 has been performed, the clock buffer 544 has been on long enough to provide a stable clock signal that may be used to synchronously transfer the data involved in the transaction.

At step 820 the core of DRAM 520 loads into an output buffer (e.g. clocked buffer 536) the data block from the column specified in the address sampled at step 814. At step 822 the data block is transmitted from the output buffer to the memory controller 518 in a synchronous fashion based on the clock signal from clock buffer 544. At step 824 it is determined if the  $\overline{\text{CAS}}$  signal went LOW again. If so, then an additional packet of data is to be sent in the current read transaction. Control therefore returns to steps 810 and 814. If the  $\overline{\text{CAS}}$  signal did not go LOW again, then the last packet of data for the transaction has been transmitted, and the clock buffer 544 is turned off at step 826.

Control proceeds to step 828 if the transaction is a write transaction. At step 828, DRAM 520 receives data through clocked buffer 536 which is driven by the clock signal from clock buffer 544. When the packet of data has been received, the packet of data is stored in the memory array 524 of DRAM 520 at step 830. At step 824 it is determined if the  $\overline{\text{CAS}}$  signal went LOW again subsequent to steps 810 and 812. If so, then an additional packet of data is to be received in the current write transaction. Due to the asynchronous control circuitry of DRAM 520, the clock does not have to be operating to perform a memory cell refresh operation. Control therefore returns to step 814. If the  $\overline{\text{CAS}}$  signal did not go LOW subsequent to step 812, then the last packet of data for the transaction has been received, and the clock buffer 544 is turned off at step 826.

It should be noted that an assertion of  $\overline{\text{RAS}}$  may be followed by any arbitrary read/write sequence. For example,  $\overline{\text{RAS}}$  may go LOW to cause a particular row of data to be loaded into sense amplifier array 560. Subsequently, a series of mixed reads and writes may be performed on the row of data. In addition, an arbitrary amount of time may elapse between  $\overline{\text{CAS}}$  signals. The duration of  $\overline{\text{RAS}}$  and the delay between CAS operations is limited only by core considerations such as refresh rates.

-20-

### MULTIPLE-DRAM SYSTEMS

The memory system 500 of Figure 5a includes only one DRAM 520. However, the present invention is not limited to memory systems with any particular number of DRAMs. Additional DRAMs may be added to memory system 500 without affecting the operations described above. Each of the DRAMs would be connected to memory controller 518 by its own private  $\overline{\text{CAS}}$  line,  $\overline{\text{RAS}}$  line and  $\overline{\text{CKE}}$  line. All of the other lines that connect the DRAMs to the memory controller 518 may be shared.

In an alternate embodiment, a memory system has a two dimensional array of memory chips. In such an embodiment, all DRAMs that belong to the same column of the two dimensional array would share the same set of control lines, while each row of DRAMs in the two dimensional array would have its own set of control lines.

### MULTIPLEXED EMBODIMENTS

In the embodiment illustrated in Figure 5a, the throughput is maximized by providing separate lines for each type of signal so that the signals which are separated can function simultaneously. However, as a general rule, the higher the number of lines required by a memory system, the more expensive it is to manufacture the components for the memory system. Therefore, the approach shown in Figure 5a may not be optimal when the cost of manufacturing is taken into account. In an alternative embodiment to that shown in Figure 5a, the number of lines is reduced by multiplexing some of the lines to allow the same lines to carry more than one type of signal.

According to an alternative embodiment, the mask signal can be sent over the address bus 504, eliminating the need for mask lines 508. Figure 9 is a timing diagram illustrating the timing of the signals generated during a write operation in such an embodiment. The timing proceeds in the same fashion as described above with reference to Figure 7 with the difference that the mask signal is not sent over a mask bus at the same time as the column address is sent over the address bus. Rather, both the address and the mask bits are sent over a combined address/mask bus, where the address bits precede the corresponding mask bits.

Similar to memory system 500, the memory controller in a combined address/mask bus embodiment indicates to the DRAM that the column address is

-21-

present on the address/mask bus by causing the  $\overline{\text{CAS}}$  signal to go LOW (at times T1 and T4). The memory controller indicates to the DRAM that the transaction is a write transaction by causing the  $\overline{\text{WE}}$  signal to be LOW at the start of the clock cycle after the rising edge of the  $\overline{\text{CKE}}$  signal. In addition, the memory controller indicates the presence of the mask bits on the address/mask bus (at times T2 and T5) by causing the  $\overline{\text{CAS}}$  signal to go HIGH. When the DRAM detects the rising edge of the  $\overline{\text{CAS}}$  signal, the DRAM reads the mask bits from the address/mask bus.

In an alternate combined address/mask bus embodiment, the memory controller indicates the presence of the mask bits on the address/mask bus by causing the  $\overline{\text{WE}}$  signal to go HIGH. When the DRAM detects the rising edge of the  $\overline{\text{WE}}$  signal, the DRAM reads the mask bits from the address/mask bus. When the rising edge of  $\overline{\text{WE}}$  is used to indicate the presence of mask bits, the mask bits for the transfer of a particular data block must be placed on the combined address/mask bus at the rising edge of  $\overline{\text{WE}}$  that corresponds to the transfer of the particular data block. In Figure 9, for example, the mask bits associated with the data block that is transferred beginning at T1 would be placed on the combined address/mask bus to be read at the first rising edge of  $\overline{\text{WE}}$  after time T0. Similarly, the mask bits associated with the data block that is transferred beginning at T4 would be placed on the combined address/mask bus to be read at the first rising edge of  $\overline{\text{WE}}$  after time T3.

For read transactions, mask bits are not transmitted. Therefore, read transactions in a combined address/mask bus embodiment proceed as illustrated in Figure 6, with the exception that the separate mask signal does not exist.

In alternative embodiments, addresses may be multiplexed even though a separate address bus is provided. For example, row addresses may be sent over the separate address bus while column addresses are multiplexed on the same bus that carries data. Similarly, column addresses may be sent over the separate address bus while row addresses are multiplexed on the same bus that carries the data.

-22-

### READ TRANSACTION IN A MULTIPLEXED DATA/ADDRESS/MASK BITS EMBODIMENT

The number of lines required by a memory system that implements the present invention may be further reduced by using the same set of lines to transmit the data, address and mask bits. Figure 10 illustrates the timing of signals generated during a read transaction in an embodiment in which the data, address and mask bits are transmitted over a combined bus.

Referring to Figure 10, at time T1 the  $\overline{\text{RAS}}$  signal goes LOW to indicate to the DRAM that a row address is on the combined bus. The DRAM reads the row address and begins a sense operation to load the appropriate row of data into the sense amplifier array. At time T2, the  $\overline{\text{CAS}}$  signal goes LOW to indicate to the DRAM that a column address is on the combined bus. In the illustrated example, the column address on the combined bus at time T2 specifies column address A. Also at time T2, the signal goes LOW to turn on the clock buffer within the DRAM.

The memory controller causes the  $\overline{\text{CKE}}$  signal to go HIGH to indicate that time T3 is the timing mark for the transfer of data from column address A. At time T4, the  $\overline{\text{CAS}}$  signal goes LOW to indicate to the DRAM that a column address is on the combined bus. In the illustrated example, the column address on the combined bus at time T4 specifies column address B. At time T5, the data from column address A begins to be placed on the combined bus. The memory controller knows to expect the data from column address A at time T5 because the time at which the data from column address A is placed on the combined bus is determined by the timing mark at time T3. Also at time T5, the DRAM core begins to access the data from column address B.

The memory controller causes the  $\overline{\text{CKE}}$  signal to go HIGH to indicate that T6 is the timing mark for the transfer of data from column address B. At time T7, the  $\overline{\text{CAS}}$  signal goes LOW to indicate to the DRAM that a column address is on the combined bus. In the illustrated example, the column address on the combined bus at time T7 specifies column address C. At time T8, the data from column address B begins to be placed on the combined bus. The memory controller knows to expect the data from column address B at time T8 because the time at which the data from column address B is placed on the combined bus is determined by the timing mark at time T6. Also at time T8, the DRAM core begins



-23-

to access the data from column address C. This process may be repeated to transfer any arbitrary number of columns of data. Each falling edge in the  $\overline{\text{CAS}}$  signal initiates a transfer that constitutes an independent transaction, and continues until the entire set of read and write transactions have been completed.

#### WRITE TRANSACTION IN A MULTIPLEXED DATA/ADDRESS/MASK BITS EMBODIMENT

Figure 11 illustrates the timing of signals generated during a write transaction in an embodiment in which the data, address and mask bits are transmitted over a combined bus. Referring to Figure 11, at time T1 the  $\overline{\text{RAS}}$  signal goes LOW to indicate to the DRAM that a row address is on the combined bus. The row decoder receives the row address and begins a sense operation to load the appropriate row of data into the sense amplifier array. At time T2, the  $\overline{\text{CKE}}$  signal goes LOW to turn on the clock buffer within the DRAM.

The memory controller causes the  $\overline{\text{CKE}}$  signal to go HIGH prior to time T3 to indicate that T3 is the timing mark. At time T3 the DRAM samples the  $\overline{\text{WE}}$  signal to determine that the transaction is a write transaction. The DRAM receives a column address specifying column A when  $\overline{\text{CAS}}$  goes LOW, and mask data when  $\overline{\text{CAS}}$  goes HIGH. The transfer of data for column A begins at time T4. The DRAM knows to receive the data for column A at time T4 because clock cycle T4 is a predetermined number of clock cycles after the timing mark (T3). In the illustrated example, data is transmitted three clock cycles after the corresponding timing mark.

The second rising edge of the  $\overline{\text{CKE}}$  signal indicates to the DRAM that time T5 is a timing mark for a second data transfer operation. The DRAM samples the  $\overline{\text{WE}}$  signal at time T5 to determine that the second data transfer transaction will be a write transaction.

#### INDEPENDENT SENSE AND PRECHARGE SIGNALS

When a single  $\overline{\text{RAS}}$  line is connected to a DRAM, only one bank within the DRAM may be sensed at any given time. Therefore, only one sense amplifier array is required per DRAM. To allow more than one bank to be sensed at a time, multiple  $\overline{\text{RAS}}$  lines can be connected to the DRAM. If each bank within the DRAM has its own  $\overline{\text{RAS}}$  line, then the controller can independently control (and

sense) each of the banks. In such an embodiment, each bank would have its own sense amplifier array. However, the cost of providing a separate line for each bank in each DRAM is significant.

To avoid the cost of providing a separate  $\overline{\text{RAS}}$  line for each bank in each DRAM, the  $\overline{\text{RAS}}$  line may be replaced with separate  $\overline{\text{SENSE}}$  and PRECHARGE signals. In this embodiment, the memory controller can cause a row within any given bank to be sensed by causing the  $\overline{\text{SENSE}}$  signal to go LOW while placing an address on the address bus that indicates a particular row and bank within the DRAM. The rising edge of  $\overline{\text{SENSE}}$  is irrelevant, though a minimum pulse width must be observed. Similarly, any bank may be precharged by causing the PRECHARGE signal to go HIGH while placing an address on the address bus that indicates a particular bank within the DRAM. The falling edge of the PRECHARGE signal is irrelevant. Figure 12 illustrates the correlation between separate  $\overline{\text{SENSE}}$  and PRECHARGE signals and a traditional  $\overline{\text{RAS}}$  signal. In this embodiment, each bank will have its own sense amplifier array, but will not require its own  $\overline{\text{RAS}}$  line.

In a system that provides separate PRECHARGE and  $\overline{\text{SENSE}}$  signals, the address that is sent when the  $\overline{\text{CAS}}$  signal goes LOW includes, in addition to a column address, bits that indicate a particular memory bank. The DRAM transmits data from the specified column of the sense amplifier array that corresponds to the specified memory bank. Thus, in the read transaction described above with reference to Figure 10, column address A, column address B and column address C may be columns in different memory banks.

In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

-25-

**CLAIMS**

What is claimed is:

1. A method for transferring data within a computer system, the method comprising the steps of:  
causing a controller to transmit control signals requesting a data transfer to a memory device;  
causing said memory device to asynchronously receive said control signals; and  
causing said memory device to synchronously perform said data transfer.
2. The method of Claim 1 wherein:  
said memory device has a first mode in which data transfer circuits within said memory device are not driven by an internal clock signal;  
said memory device has a second mode in which data transfer circuits within said memory device are driven by said internal clock signal;  
the method further includes the step of causing said memory device to assume said second mode in response to one or more of said control signals; and  
the step of causing said memory device to synchronously perform said data transfer includes causing said memory device to transfer data with said data transfer circuits while said data transfer circuits are being driven by said internal clock signal.
3. The method of Claim 2 wherein:  
said memory device includes a timer;  
the method further comprises the steps of  
loading said timer with a value in response to at least one of said control signals; and  
causing said memory device to assume said first mode when said timer reaches a predetermined value.
4. The method of Claim 3 further including the step of updating said timer while said memory device is in said second mode in response to one or more control signals received by said memory device subsequent to said control signals.
5. The method of Claim 2 further comprising the steps of:

-26-

causing said controller to send a control signal to said memory device after said memory device has assumed said second mode;  
causing said memory device to select a clock cycle of said internal clock signal as a timing mark based on said control signal; and  
causing said memory device to time said data transfer based on said timing mark.

6. The method of Claim 5 wherein:  
the step of causing said memory device to assume said second mode in response to one or more of said control signals comprises the step of causing said memory device to assume said second mode in response to a first edge of a clock enable signal; and  
the step of causing said memory device to select a clock cycle of said internal clock as a timing mark based on said control signal comprises the step of causing said memory device to select a clock cycle of said internal clock as a timing mark based on a second edge of said clock enable signal.
7. The method of Claim 1 wherein:  
the step of causing a controller to transmit control signals requesting a data transfer to a memory device comprises the step of causing said controller to transmit an address over a first set of lines, said address indicating a location of data which will be involved in said data transfer; and  
the step of causing said memory device to synchronously perform said data transfer includes the step of causing data to be transferred over said first set of lines.
8. The method of Claim 2 wherein:  
the step of causing a controller to transmit control signals requesting a data transfer to a memory device comprises the step of causing said controller to transmit an address over a first set of lines, said address indicating a location of data which will be involved in said data transfer; and  
the step of causing said memory device to synchronously perform said data transfer includes the step of causing data to be transferred over said first set of lines.

-27-

9. The method of Claim 2 wherein:  
the step of causing a controller to transmit control signals includes causing said controller to transmit control signals that specify a read transaction;  
the memory device performs a core access in response to said control signals; and  
the step of causing said memory device to assume said second mode is performed while the memory device is performing the core access.
10. The method of Claim 9 wherein the step of causing said memory device to assume said second mode includes the step of starting to generate said internal clock signal.
11. The method of Claim 7 wherein the step of causing said controller to transmit an address over a first set of lines includes the step of causing said controller to transmit a column address over said first set of lines.
12. The method of Claim 7 wherein the step of causing said controller to transmit an address over a first set of lines includes the step of causing said controller to transmit a row address over said first set of lines.
13. The method of Claim 2 wherein:  
the memory device includes a clock generation circuit that has a first configuration and a second configuration;  
the method includes the step of generating said internal clock by passing an external clock signal through a delay lock loop circuit when the clock generation circuit is in said first configuration; and  
the method includes the step of generating said internal clock without passing the external clock signal through a delay lock loop circuit when the clock generation circuit is in the second configuration.
14. The method of Claim 13 further including the steps of:  
generating the external clock signal at a first frequency when the clock generation circuit is in the first configuration; and  
generating the external clock signal at a second frequency when the clock generation circuit is in the second configuration, wherein the first frequency is higher than said second frequency.
15. The method of Claim 1 wherein:

-28-

the memory device includes a clock generation circuit that has a first configuration and a second configuration;

when the clock generation circuit is in the first configuration, the step of causing said memory device to synchronously perform said data transfer includes the step of driving data transfer circuits within said memory device by an internal clock signal generated by passing an external clock signal through a delay lock loop circuit; and

when the clock generation circuit is in the second configuration, the step of causing said memory device to synchronously perform said data transfer includes the step of driving data transfer circuits within said memory device by an internal clock signal generated based on said external clock signal without passing said external clock signal through said delay lock loop circuit.

16. The method of Claim 15 further comprising the steps of:  
generating said external clock signal at a first frequency when the clock generation circuit is in the first configuration; and  
generating said external clock signal at a second frequency when the clock generation circuit is in the second configuration, wherein the first frequency is higher than the second frequency.
17. The method of Claim 15 wherein:  
the clock generation circuit has a third configuration; and  
said external clock signal does not drive said data transfer circuits when said clock generation circuit is in said third configuration;  
the method further comprising the steps of  
causing said controller to transmit a second set of control signals requesting a second data transfer;  
causing said memory device to asynchronously receive said second set of control signals; and  
causing said memory device to asynchronously perform said second data transfer while said clock generation circuit is in said third configuration.

-29-

18. The method of Claim 17 further comprising the step of driving said data transfer circuits with a control signal other than said external clock signal when said clock generation circuit is in said third configuration.
19. The method of Claim 1 further comprising the steps of:
  - causing said controller to transmit a second set of control signals requesting a second data transfer;
  - causing said memory device to asynchronously receive said second set of control signals; and
  - causing said memory device to asynchronously perform said second data transfer.
20. The method of Claim 1 wherein:
  - the control signals include a first control signal and a second control signal;
  - the memory device begins generation of an internal clock signal in response to said first control signal;
  - the memory device takes a sample of the second control signal at a time that is based on when an edge of said first control signal occurs; and
  - the memory device determines whether the data transfer will be a write transaction or a read transaction based on the sample.
21. The method of Claim 20 wherein:
  - the memory device begins generation of the internal clock signal in response to detecting a particular edge of said first control signal; and
  - the memory device takes the sample of the second control signal in response to detecting said particular edge of said first control signal.
22. The method of Claim 2 further including the steps of:
  - the memory device detecting whether any outstanding transactions for said memory device exist; and
  - if no outstanding transactions for said memory device exist, then causing said memory device to assume said first mode.
23. The method of Claim 1 wherein the step of causing said memory device to synchronously perform said data transfer includes causing said memory device to transfer data that corresponds to a plurality of column addresses between said memory device and said controller over a data bus.

-30-

24. The method of Claim 23 wherein said data transfer includes both writing data to said memory device and transmitting data from said memory device.
25. The method of Claim 1 wherein:
  - the step of causing said controller to transmit control signals includes causing said controller to transmit an address and mask information;
  - the address identifies a storage location within the memory device that will be involved in the data transfer;
  - the address and the mask information are sent to the memory device over the same control lines.
26. The method of Claim 1 wherein:
  - the step of causing said controller to transmit control signals includes causing said controller to transmit mask information;
  - the step of causing said memory device to synchronously perform said data transfer includes causing said memory device to store data sent to said memory device over a bus; and
  - the controller sends the mask information to the memory device over said bus.
27. The method of Claim 26 wherein:
  - the step of causing said controller to transmit control signals further includes causing said controller to transmit address information;
  - the address information indicates a storage location within the memory device in which the data is to be stored; and
  - the controller sends the address information to the memory device over said bus.
28. A memory device for storing information, comprising:
  - a memory array;
  - a sense amplifier array coupled to said memory array;
  - decode circuitry coupled to said sense amplifier array, said decode circuitry causing information to be transferred from said memory array to said sense amplifier array in response to asynchronously received control signals;



-31-

- a storage circuit for temporarily storing information, said storage unit having a clock input, said storage circuit being coupled between said amplifier array and a data bus that is external to said memory device; and
- said storage circuit with a clock input being configured to synchronously transmit said information to a device external to said memory device when said clock input is driven by a clock signal.
29. The memory device of Claim 28 wherein the decode circuitry includes:
- a row decoder coupled to said memory array;
  - said row decoder causing information to be transferred from said memory array to said sense amplifier array in response to a first address;
  - a column decoder coupled to said sense amplifier array; and
  - said column decoder causing information to be transferred from said sense amplifier array through said output circuits in response to a second address.
30. The memory device of Claim 28 wherein:
- the memory device further comprises a clock buffer having an signal input, a control input and a signal output;
  - the signal input of said clock buffer is coupled to an external clock line;
  - the signal output of said clock buffer is coupled to the clock input of said storage circuit;
  - the clock buffer is configured to assume one of a first mode and a second mode in response to a control signal applied to said control input;
  - any clock signal applied at said signal input of said clock buffer is not transmitted to said signal output when said clock buffer is in said first mode;
  - any clock signal applied at said signal input of said clock buffer is transmitted to said signal output when said clock buffer is in said second mode.
31. The memory device of Claim 30 wherein
- the memory device further includes control circuitry coupled to apply said control signal to the control input of said clock buffer;
  - said control circuitry is coupled to a line that is external to said memory device;

-32-

said control circuitry causes said clock buffer to assume said second mode in response to a particular signal on said line that is external to said memory device; and

said control circuitry causes said clock buffer to assume said first mode after said information is synchronously transmitted from said buffer.

32. The memory device of Claim 28 wherein the memory device asynchronously receives said control signals over said data bus.
33. A memory system comprising:
  - a memory device having a plurality of modes including a first mode and a second mode;
  - the memory device assuming a lower power state in which data transfer circuits within said memory device are not driven by an internal clock signal when said memory device is in said first mode;
  - the memory device assuming a higher power state in which said data transfer circuits are driven by said internal clock signal when said memory device is in said second mode;
  - a memory controller coupled to said memory device by a plurality of lines;
  - the memory controller asynchronously sending a first set of control signals that specify a data transfer transaction to said memory device;
  - the memory controller asynchronously transmitting a clock control signal to said memory device to cause said memory device to assume said second mode; and
  - the memory device assuming said second mode in response to said clock control signal and driving said data transfer circuits with said internal clock signal to perform the data transfer transaction specified by said first set of control signals.
34. The memory system of Claim 33 wherein:
  - the internal clock signal is generated from an external clock signal;
  - the memory controller sends a control signal to identify a particular clock cycle of said external clock signal after said memory device has assumed said second mode and prior to performance of said data transfer transaction; and

-33-


the memory device times performance of said data transfer transaction based on said particular clock cycle.

35. The memory system of Claim 33 wherein:  
said memory device includes a plurality of memory banks;  
said memory device includes a sense amplifier array for each of said plurality of memory banks;  
said memory controller sends to said memory device a sense control signal and an address that indicates a particular row of a particular memory bank within said memory controller;  
in response to said sense control signal and said address, said memory device causes information stored in said particular row of said particular memory bank to be sensed by the sense amplifier array that corresponds to said particular memory bank;  
said memory controller sends to said memory device a precharge control signal and an address that indicates a particular memory bank within said memory controller, wherein said precharge control signal is sent on a separate line from said sense control signal; and  
in response to said precharge control signal and said address, said memory device causes said particular memory bank to be precharged.
36. A method for performing a data transfer in a computer system, the method comprising the steps of:  
causing a memory device having data transfer circuitry to assume a first mode in which said data transfer circuitry is not driven by any clock signal;  
sending control information and address information to said memory device, said control information specifying a data transfer transaction, said address information specifying storage within said memory device that will be involved in said data transfer transaction;  
in response to said control information, causing said memory device to assume a second mode in which said data transfer circuitry is driven by a clock signal; and


-34-

causing said memory device to perform said data transfer transaction in synchronism with said clock signal while said memory device is in said second mode.


37. The method of Claim 36 further comprising the steps of:  
sending second control information and second address information to said memory device, said second control information specifying a second data transfer transaction, said second address information, specifying storage within said memory device that will be involved in said second data transfer transaction; and  
causing said memory device to perform said second data transfer transaction asynchronously while said memory device is in said first mode.
38. The method of Claim 36 further comprising the step of causing said memory device to assume said first mode after performing said data transfer transaction.
39. The method of Claim 38 wherein the step of causing said memory device to assume said first mode after performing said data transfer transaction includes the steps of:  
detecting if a predetermined amount of time has elapsed since receipt by the memory device of a particular control signal; and  
causing said memory device to assume said first mode when said predetermined amount of time has elapsed since receipt by the memory device of said particular control signal.
40. The method of Claim 39 wherein:  
the method further includes the step of updating a value in a timer every time the memory device receives said particular control signal; and  
the step of determining if a predetermined amount of time has elapsed since receipt by the memory device of a particular control signal includes the step of determining if said value in said timer has reached a predetermined value.
41. The method of Claim 36 wherein the step of sending control information to said memory device includes causing a memory controller to asynchronously transmit said control information to said memory device.




**FIG. 1**  
**(PRIOR ART)**



**FIG. 2**  
**(PRIOR ART)**



**FIG. 3**  
(PRIOR ART)



**FIG. 4**  
(PRIOR ART)



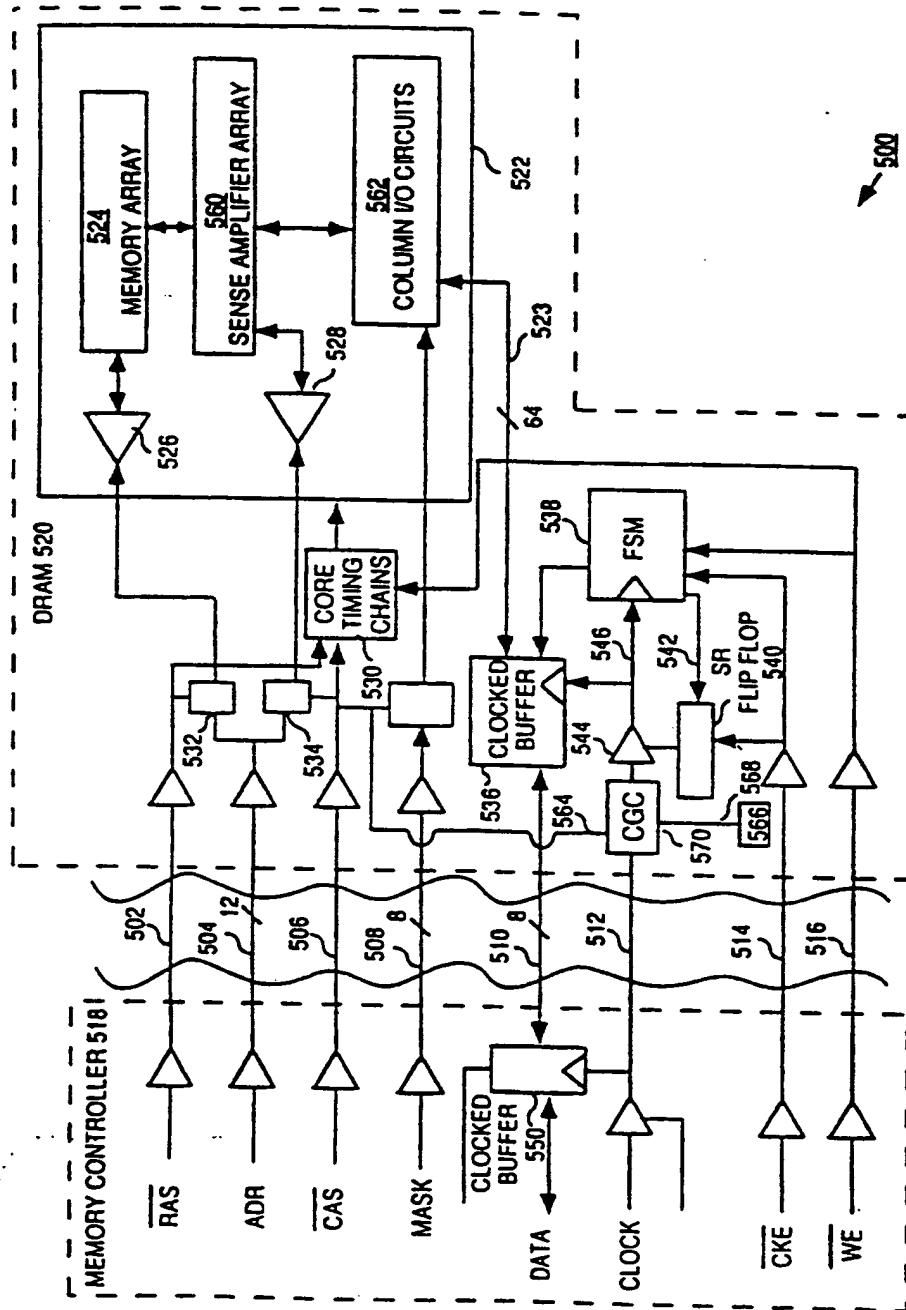



FIG. 5a



**FIG. 5b**




FIG. 6

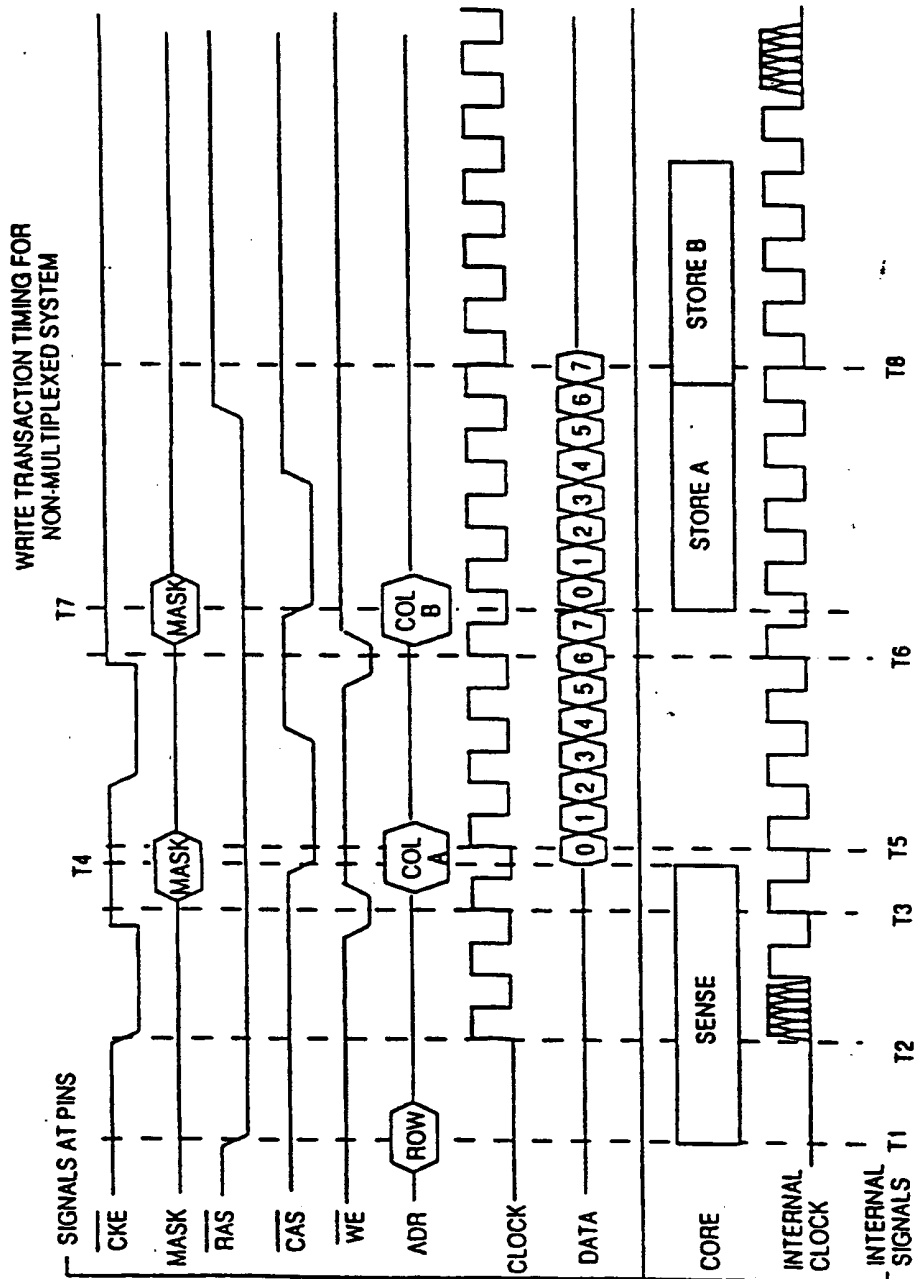


Fig. 7

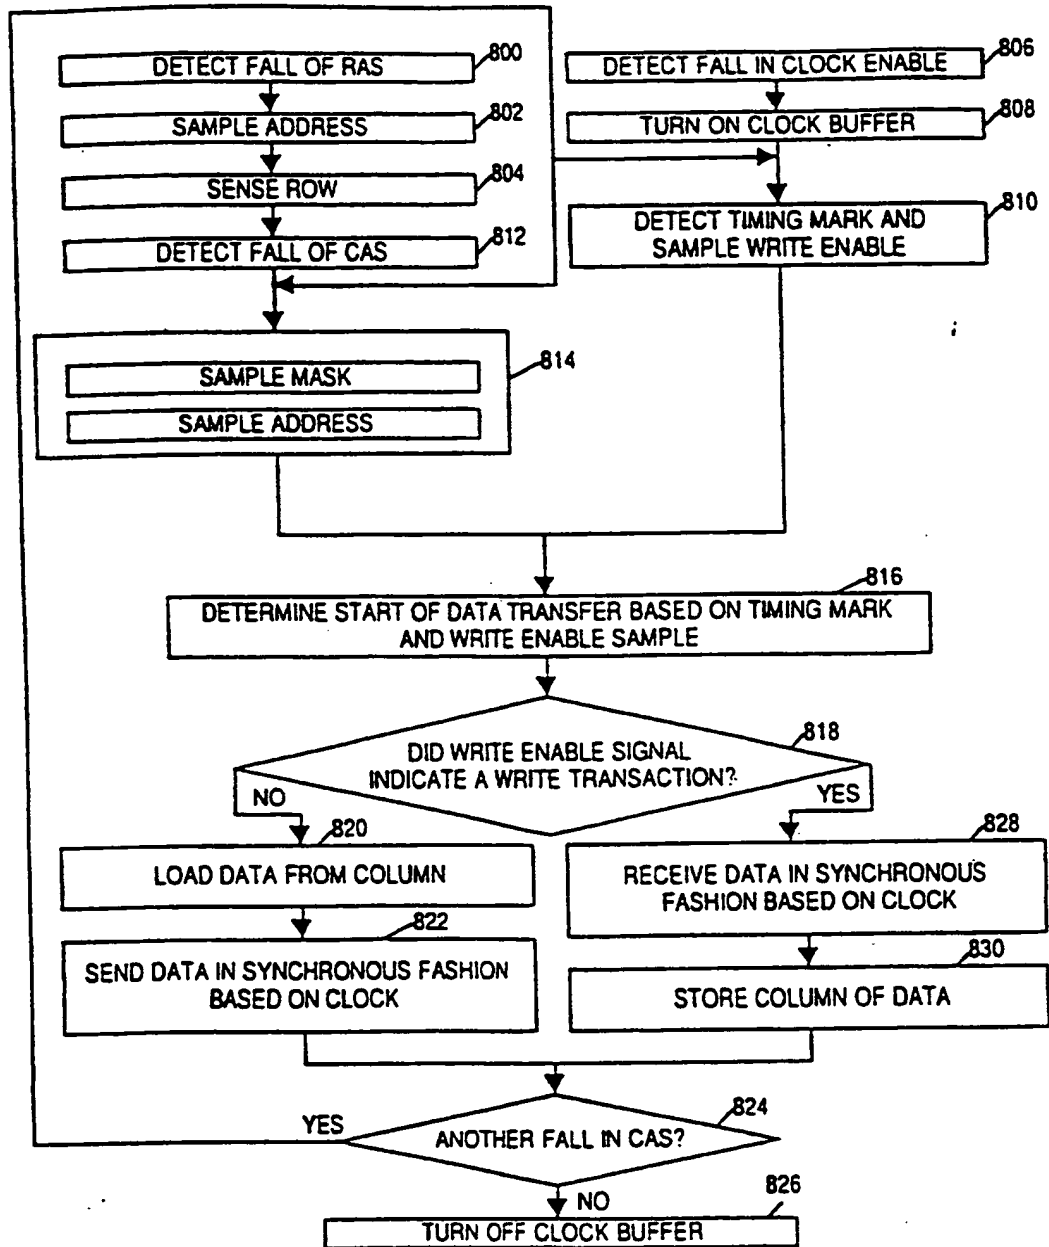


FIG. 8




FIG. 9




Fig. 10





Fig. 11





SEPARATE SENSE AND  
PRECHARGE CONTROL SIGNALS

Fig. 12

**THIS PAGE BLANK (USPTO)**